

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 220 920 B1

(12)

EUROPEAN PATENT SPECIFICATION

- (43) Date of publication of patent specification: 24.07.91 (51) Int. Cl.⁵: G06F 9/46, G06F 12/14
(21) Application number: 86308166.7
(22) Date of filing: 21.10.86

(54) Instruction for implementing a secure computer system.

(30) Priority: 28.10.85 US 792702

(43) Date of publication of application:
06.05.87 Bulletin 87/19

(45) Publication of the grant of the patent:
24.07.91 Bulletin 91/30

(84) Designated Contracting States:
CH DE FR GB IT LI NL SE

(56) References cited:
US-A- 4 177 510

COMPUTER DESIGN, vol. 20, no. 1, January
1981, pages 111-120, Winchester, Mass., US;
S. WALLACH et al.: "32-bit minicomputer
achieves full 16-bit compatibility"

(73) Proprietor: Hewlett-Packard Company
Mail Stop 20 B-O, 3000 Hanover Street
Palo Alto, California 94304(US)

(72) Inventor: Miller, Terrence C.
169 Oak Court
Menlo Park, CA 94025(US)
Inventor: Baum, Allen
2310 Cornell Street
Palo Alto, CA 94306(US)
Inventor: Bryg, William R.
18630 Perego Way
Saratoga, CA 95070(US)
Inventor: Mahon, Michael J.
1975 Adele Place
San Jose, CA 95125(US)

(74) Representative: Colgan, Stephen James et al
CARPMAELS & RANSFORD 43 Bloomsbury
Square
London WC1A 2RA(GB)

EP 0 220 920 B1

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid (Art. 99(1) European patent convention).

Description

In hierarchical computer systems, protection checking must occur whenever a process having lower privilege desires to call on a system wide service routine having higher privilege. Several levels of privilege are usually provided in such systems including a highest privilege available only to the operating system, one or more intermediate privilege levels available to either the operating system or intermediate program supervisors, and a lowest privilege level available to unprivileged operations and to user programs. When a user program needs to call the computer operating system to do something which can only be accomplished by the operating system itself, such as access a section of the physical memory in a virtually addressed computer memory, some form of protection must be provided so that the user program can only use the highly privileged routine in an orderly way. Without some form of protection there would be no data security in the system and any program would have unlimited access to all of the system data or a user program could actually destroy the operating system itself, both of which are unacceptable in most commercial systems. For example, if a payroll file is stored on the computer system and the system lacked privilege protection, someone could merely write a user program to give himself a raise.

Such protection checking is a frequent operation that usually involves significant system resources in validating the privileges and access rights of the calling routine, resulting in a substantial degradation in system performance. Previous systems have frequently used "supervisor call" instructions to limit the number of entry points available into the operating system. With such supervisor call instructions, a user program seeking higher privilege actually causes a hardware interrupt of the system processor, during which the system will branch to a single operating system location. The operating system must then figure out what it was that the user wanted to do, branch out through a large privilege dispatch table, perform the desired operation, return to the single operating system location, store all the desired system states, and return to the user program. Not only is such a supervisor call time consuming, but such a system requires the supervisor call to be a unique instruction which is different from other normal procedure calls.

US-A-4177510 describes a system for preventing unauthorised access from a lower privilege level to a higher privilege level in a hierarchical computer system. However, not only is the calling checked to see if there is sufficient access rights to make a procedure call, but also a second check is

made to ensure that the location of code to be executed is within an approved list. Therefore, two checks are required before the higher privilege is granted.

The present invention uses a "gateway" instruction which is a branch instruction enabling privilege checking in a computer system with only one check on user access. In addition, the gateway instruction permits an unlimited number of entry points into the operating system. Higher privilege services can be called using the same subroutine calling convention as is used for calling other procedures within the computer system so that the code that is compiled in the lower privileged routines in seeking a higher privilege level is the same as a "normal" procedure call. If the page of virtual memory on which the called entry point resides is one of a set of gateway types, then the gateway instruction will cause the routine's privilege to be raised to the level specified by the page itself, permitting the further desired execution requiring the higher privilege level, unless the routine's privilege is already greater than that of the page, in which case no change in privilege level is required.

The actual security check is performed by first checking the state of the privilege level of the currently running process against the bits in a seven bit access rights field associated with the virtual page address of the called routine during the virtual address translation performed by a Translation Lookaside Buffer (TLB). If the caller's privilege level is within bounds, the TLB permits the calling routine to execute instructions on the targeted virtual page. Execution of the gateway instruction will then proceed to branch to the location indicated by the gateway with the needed privilege level supplied by other bits in the access rights field of the TLB entry. However, if a calling routine attempts to execute on a page which requires higher or lower privilege than the routine currently has, the TLB will inhibit execution of instructions, including gateway instructions, with a software trap, so that the calling routine never performs the unpermitted instruction. Since this protection scheme can be nearly the same as that used to prevent the performance of other instructions from an unpermitted page, the gateway protection can be structurally the same as other instructions in the system.

Thus, the gateway instruction acts like a normal branch instruction relative to the state of the program counter (i.e., a PC relative branch). In addition, the gateway instruction also saves the original privilege level of the calling routine in a general register specified by a field of the gateway branch. This register is, by convention, the register containing the return link to the calling routine, the low-order bits of which are the two-bit privilege level of the caller. Since execution of the gateway

instruction forces these two bits to the caller's actual privilege level, it is impossible for the calling routine to forge a return link which could cause return to the caller with higher privilege.

Also, if the routine called by the calling routine requires the services of an even more privileged routine, the called routine may itself call through another gateway, as before.

Finally, in a computer system which permits delayed branches, such as in a conventional pipelined machine in which delayed instructions are not actually executed until a given number of machine cycles (the delay period) after the delayed instruction is begun, one caution must be exercised with the gateway instructions: The calling routine must not be permitted to execute a gateway instruction in the delay period of a taken branch, since this would result in execution of the gateway (and privilege promotion) followed by a change of control back to the calling routine at higher privilege. The effect of such an operation would be for the calling routine to effectively promote its own privilege level without control of the operating system, thus destroying the entire protection checking mechanism.

Brief Description of the Drawing

Figure 1 shows an overall computer system block diagram for utilizing a protection system according to a preferred embodiment of the present invention.

Figure 2 shows a flow chart for protecting the system shown in Figure 1.

Figure 3 shows the structure of a Target Register for holding the original privilege level of a calling routine during the execution of a gateway instruction according to the present invention.

Figure 4 shows a block diagram of a memory page containing a gateway instruction according to a preferred embodiment of the present invention.

Figure 5 shows a table containing access rights for use according to a preferred embodiment of the present invention.

Figure 6 shows a processor status word with the system shown in Figure 1.

Description of the Preferred Embodiment

Figure 1 shows a block diagram of a pipelined computer system 10 for using a gateway instruction according to the present invention and Figure 2 shows a flow chart for using the gateway instruction for protecting the system 10. An Instruction Unit 20 contains a low privilege routine which requires a procedure call to a higher privileged service routine. The Instruction Unit 20 seeks this higher privileged routine by addressing the Trans-

lation Lookaside Buffer (TLB) 30 via the Virtual Address Bus 35 to determine the location in Physical Memory 40 containing an appropriate entry point of a gateway instruction. Typically, the various entry points of the gateway instructions are published within the system documentation for programming use. The TLB 30 calculates the address of the desired entry point within the Physical Memory 40 and a gateway instruction located at the calculated address is then transmitted via a Next Instruction Bus 50 from the Physical Memory 40 to the Instruction Unit 20, to an Execution Unit 60 and to a physical Target Register 70 within the Register File 80.

A return address for returning from the higher privileged service routine is then stored in the Target Register 70 by the Instruction Unit 20 via a Results Bus 85. The Target Register 70 as shown in Figure 3 contains the return address in Address Location 300 with the original, lower privilege level stored in two lower order bits 310. The TLB 30 then checks the access rights of the calling instruction as will be described shortly to determine if execute access is permitted. If execute access is denied by the TLB 30, a software trap is transmitted from the TLB 30 to the Instruction Unit 20 on a Trap Control Bus 90 to halt execution of the gateway instruction in the Execution Unit 60. If execute access is allowed by the TLB 30, and no delayed taken branch is pending, the gateway instruction resaves the actual privilege level of the calling routine in the two low-order bits of Target Register 70 (to rule out forgery by the calling routine), and raises the privilege level of the calling routine to the privilege level specified within the page type field 412 of the TLB entry for the page containing the gateway instruction, and a target address for branching to a called routine is calculated in either the Instruction Unit 20 or the Execution Unit 60, as appropriate. A target instruction located at the target address is then fetched from the physical Memory 40 on the next instruction cycle of the system 10 for use in the Instruction Unit 20 and execution of the called service routine having the desired higher privilege proceeds in the Execution Unit 60. The gateway instruction therefore performs as a delayed branch from the entry point to the target instruction.

After the finally called service routine is completed, the Execution Unit 60 reads the return address stored in the Target Register 70 via the Results Bus 85 and returns to the calling routine at the specified return address with the original lower privilege stored in the Target Register 70.

The actual security protection by the TLB 30 is the same for "normal" instructions as for accessing a gateway instruction, and is performed with a granularity of an entire page of virtual memory.

That is, once the execution access rights of a calling routine have been verified by the TLB 30, the calling routine will have execute access to all of the information on a virtual memory page 410, including a gateway instruction 430 as shown in Figure 4. Each virtual page 410 has associated with it a page characteristic 400 containing access rights to the page and an access identifier. The page characteristic 400 is compared to the calling routine's access type (read, write, or execute), privilege level, and a set of protection identifiers to check if the calling routine is allowed to read, write or execute on that memory page 410. The access rights are encoded in the seven bits of the page characteristic 400 with the first three bits 412 specifying a page type, and the last four bits 415 being two privilege level fields specifying the most or least privilege that the calling routine can have to be permitted to use the page 410. Figure 5 shows how the access rights are encoded in the page characteristic 400, with XLEAST and XMOST being the least and most privileged levels that can execute, and READ and WRITE fields specifying the least privileged field that can have access to the page 410. Thus there are three types of access possible (read, write, and execute) for "normal" system functions, with loads checking only read access, stores checking only write access, and instruction fetch checking only execute access within XLEAST to XMOST bounds. In general, four different types of gateway pages as shown in Figure 5 (i.e., Proprietary/Gateways 0, 1, 2, and 3) are sufficient for defining the desired number of different privilege levels within the system 10. With four different privilege levels, only two of bits 412 are required to define the four different states (00, 01, 10, and 11) of the gateways, and it is these two bits which determine the privilege level resulting from execution of a gateway instruction on the corresponding page.

As mentioned previously, gateway instructions cannot be permitted during the pipeline delay period of a taken branch instruction. In order to prevent this occurrence, a B-bit in a status word 600 as shown in Figure 6 within the Instruction Unit 20 (see Figure 1) is set to indicate whether a taken delay branch is pending. The B-bit is set by the Instruction Unit 20 on any taken branch, is true during the pipeline delay period, and is cleared by the Instruction Unit 20 on the next pipeline cycle after the delay period. Thus, as shown in Figure 2, if a gateway instruction is attempted while a taken branch is pending, the gateway will be trapped as invalid to prevent the calling routine from promoting its own privilege level.

Claims

1. A method of preventing unauthorised access from a lower privilege level to a higher privilege level in a hierarchical computer system, said computer system having a memory and an instruction unit, which method is carried out when a calling routine having an original low privilege level desires to call on a service routine having a higher privilege level than the calling routine is normally permitted to access, the method comprising the steps of reading an access rights field contained in the instruction unit to determine current access rights of the calling routine, reading an access rights field of a page in the memory containing a gateway instruction indicated by the calling routine, and comparing the access rights field of the calling routine to the access rights field of the page in memory containing the gateway instruction indicated by the calling routine; the method being characterised by:
 - (a) raising the low privilege level of the calling routine to the higher privilege level specified by the page containing the gateway instruction if the access rights field of the calling routine indicates that the calling routine is permitted entry to the page containing the gateway instruction;
 - (b) storing the low privilege level of the calling routine in a physical target register of the calling routine, so that the calling routine cannot forge its privilege level;
 - (c) branching to a location of the service routine as specified by the gateway instruction under control of the service routine;
 - (d) executing the service routine; and
 - (e) returning control to the calling routine at an address specified by the target register with the original low privilege level stored in the target register.
2. A method according to claim 1 further comprising the step of trapping the execution of the gateway instruction if the access rights field of the calling routine does not indicate that the calling routine is permitted entry to the page containing the gateway instruction.
3. A method according to claim 1 wherein the computer system has delayed branch instructions with a specified delay period, comprising the step of trapping the execution of the gateway instruction if the delayed branch is pending when the access rights of the calling routine and the access rights of the page containing the gateway instruction are compared.

Revendications

1. Un procédé destiné à empêcher, dans un système d'ordinateur hiérarchique, un accès non autorisé depuis un niveau de privilège inférieur vers un niveau de privilège supérieur, ledit système d'ordinateur comportant une mémoire et une unité d'instructions ledit procédé étant exécuté lorsqu'un programme d'appel d'un bas niveau de privilège à l'origine souhaite appeler un programme de service à niveau de privilège supérieur à celui auquel le programme appelant est normalement autorisé à accéder, le procédé comprenant les étapes consistant à lire un champ de droits d'accès contenu dans l'unité d'instructions pour déterminer les droits d'accès actuels du programme appelant, à lire un champ de droits d'accès d'une page dans la mémoire contenant une instruction de passage indiquée par le programme appelant, et à comparer le champ de droits d'accès du programme appelant avec le champ de droits d'accès de la page de mémoire contenant l'instruction de passage indiquée par le programme appelant; le procédé étant caractérisé par les étapes consistant à:
 - (a) élever le bas niveau de privilège du programme appelant au niveau supérieur de privilège spécifié par la page contenant l'instruction de passage si le champ de droits d'accès du programme appelant indique que le programme appelant est autorisé à entrer dans la page contenant l'instruction de passage;
 - (b) mémoriser le bas niveau de privilège du programme appelant dans un registre de cible physique du programme appelant, de façon que le programme appelant ne puisse pas falsifier son niveau de privilège;
 - (c) effectuer un branchement vers un emplacement du programme de service tel que spécifié par l'instruction de passage sous la commande du programme de service;
 - (d) exécuter le programme de service; et
 - (e) renvoyer la commande au programme appelant, à une adresse spécifiée par le registre de cible, avec le bas niveau de privilège d'origine mémorisé dans le registre de cible.
2. Un procédé selon la revendication 1 comprenant en outre l'étape consistant à piéger l'exécution de l'instruction de passage si le champ de droits d'accès du programme appelant n'indique pas que le programme appelant est autorisé à entrer dans la page contenant l'instruction de passage.
3. Un procédé selon la revendication 1 dans le-

quel le système d'ordinateur possède des instructions de branchement retardées d'une période de retard spécifiée, comprenant l'étape consistant à piéger l'exécution de l'instruction de passage si le branchement retardé est en cours lorsque les droits d'accès du programme appelant et les droits d'accès de la page contenant l'instruction de passage sont comparés.

Patentansprüche

1. Verfahren zum Vorbeugen gegen unautorisierten Zugriff von einer niedriger privilegierten Stufe auf eine höher privilegierte Stufe in einem hierarchischen Computersystem, wobei das Computersystem einen Speicher und eine Befehlseinheit umfaßt, wobei das Verfahren ausgeführt wird, wenn eine anfordernde Routine einer ursprünglich gering privilegierten Stufe eine Dienstroutine aufrufen will, welche eine Stufe hat, die höher privilegiert ist, als daß die anfordernde Routine normalerweise darauf zugreifen dürfte, wobei das Verfahren die folgenden Schritte umfaßt: Auslesen eines Zugriffsrechtsfeldes in der Befehlseinheit, um aktuelle Zugriffsrechte der anfordernden Routine festzulegen, Lesen eines Zugriffsrechtsfeldes auf einer Seite in dem Speicher, die einen von der anfordernden Routine angezeigten Übergabebefehl beinhaltet, und Vergleichen des Zugriffsrechtsfeldes der anfordernden Routine mit dem Zugriffsrechtsfeld der Seite in dem Speicher, die den von der anfordernden Routine angezeigten Übergabebefehl beinhaltet, wobei das Verfahren **gekennzeichnet** ist durch:
 - a) Anheben der gering privilegierten Stufe der anfordernden Routine auf die höher privilegierte Stufe, die von der den Übergabebefehl beinhaltenden Seite bestimmt ist, wenn das Zugriffsrechtsfeld der anfordernden Routine anzeigt, daß der anfordernden Routine der Zugang zu der den Übergabebefehl beinhaltenden Seite gewährt wird;
 - b) Speichern der gering privilegierten Stufe der anfordernden Routine in einem physikalischen Zielregister der anfordernden Routine, so daß die anfordernde Routine ihre Privilegierungsstufe nicht verfälschen kann;
 - c) Springen zu einer Stelle der Dienstroutine, wie sie von dem Übergabebefehl unter Steuerung durch die Dienstroutine bestimmt ist;
 - d) Ausführen der Dienstroutine; und
 - e) Zurückgeben der Steuerung an die anfordernde Routine an einer Adresse, die von dem Zielregister bestimmt ist, und zwar mit der ursprünglichen gering privilegierten Stufe.

fe, die in dem Zielregister gespeichert ist.

2. Verfahren nach Anspruch 1, das ferner folgenden Schritt umfaßt: Stoppen der Ausführung des Übergabebefehls, wenn das Zugriffsrechtsfeld der anfordernden Routine nicht anzeigt, daß der anfordernden Routine Zugang zu der den Übergabebefehl beinhaltenden Seite gewährt wird.

5

10

3. Verfahren nach Anspruch 1, wobei das Computersystem verzögerte Sprungbefehle mit einer bestimmten Verzögerungszeit umfaßt, die den Schritt des Stoppens der Ausführung des Übergabebefehls beinhalten, wenn der verzögerte Sprung ansteht, wenn die Zugriffsrechte der anfordernden Routine und die Zugriffsrechte der den Übergabebefehl enthaltenden Seite verglichen werden.

15

20

25

30

35

40

45

50

55




FIG 1

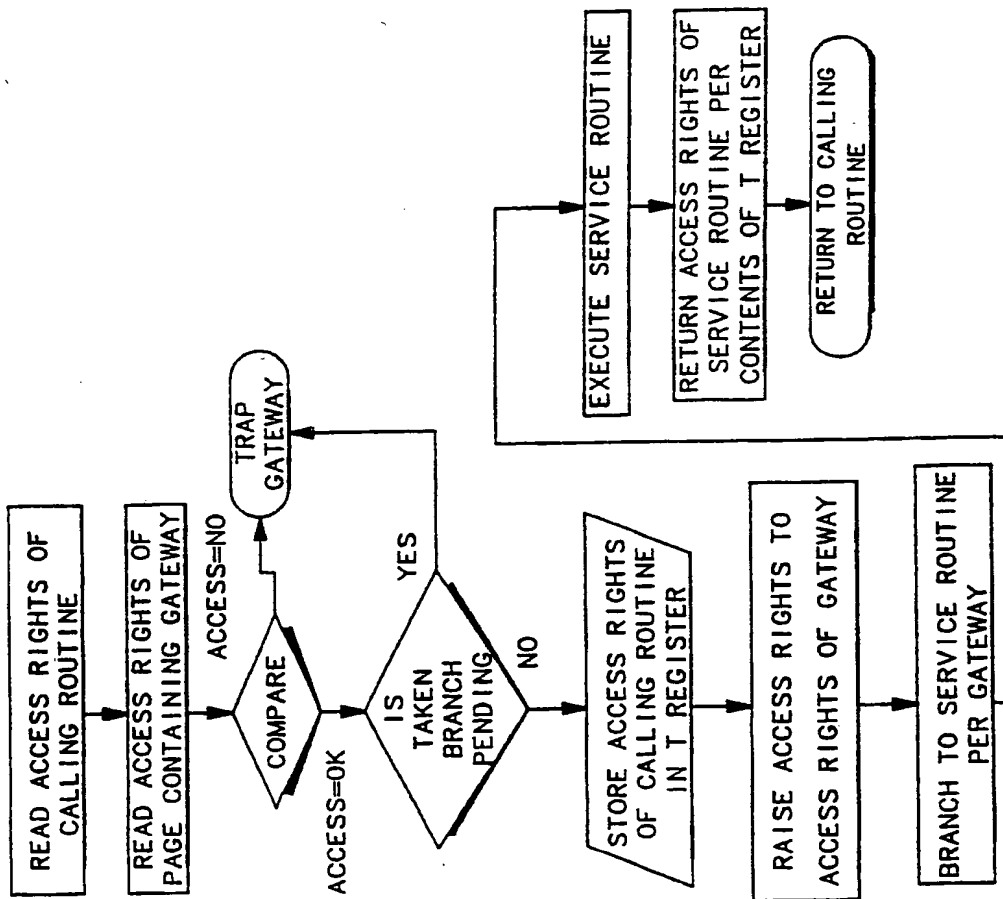





FIG 2



TYPE	PRIVILEGE LEVEL 1	PRIVILEGE LEVEL 2	USE	PRIVILEGE LEVEL AFTER EXECUTION OF GATEWAY
000	READ	—	READ ONLY DATE	
001	READ	WRITE	NORMAL DATA	
010	READ/XLEAST	XMOST	NORMAL CODE	
011	READ/XLEAST	WRITE/XMOST	DYNAMIC CODE	
100	XLEAST	XMOST	PROPRIETARY/GATEWAY	0
101	XLEAST	XMOST	PROPRIETARY/GATEWAY	1
110	XLEAST	XMOST	PROPRIETARY/GATEWAY	2
111	XLEAST	XMOST	PROPRIETARY/GATEWAY	3

FIG 5